TECHNOLOGY BRIEF

November 1995

Compaq Computer Corporation

CONTENTS

Introduction **3**

1000+ Users in an OLTP Environment **3**
OLTP Benchmark Description . 4
Results 4
Interpretation of Results 5

1000+ Users in the File Services Environment **5**

File Services Benchmark Description **6**
Test Setup 7
Results 8
Interpretation of Results 9
Extrapolation of NetBench Results to Typical Usage 10

Conclusion **11**

References **12**

Compaq ProLiant 4500 Servers in 1000+ User Environments

This paper presents results from a continuing performance investigation being conducted on Compaq servers. These results clearly indicate that the Compaq ProLiant 4500 is fully capable of supporting more than 1000 users in both on-line transaction processing and file services environments. This claim is based on analysis of industry standard benchmark results. For each of the cases studied there is:

- *an introduction to the "real world" environment simulated by the benchmark*
- *an overview of the benchmark used*
- *and an interpretation of the results obtained during the benchmark, including extrapolation to real world environments*

The paper concludes with a brief summary of the results.

It should be noted that the results presented here are from UnixWare and NetWare environments. Compaq research and experience indicates that the results can be extrapolated to other environments.

COMPAQ

Downloaded from www.Manualslib.com manuals search engine

NOTICE

The information in this publication is subject to change without notice.

COMPAQ COMPUTER CORPORATION SHALL NOT BE LIABLE FOR TECHNICAL OR EDITORIAL ERRORS OR OMISSIONS CONTAINED HEREIN, NOR FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES RESULTING FROM THE FURNISHING, PERFORMANCE, OR USE OF THIS MATERIAL.

This publication does not constitute an endorsement of the product or products that were tested. The configuration or configurations tested or described may or may not be the only available solution. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal, state or local requirements. Compaq does not warrant products other than its own strictly as stated in Compaq product warranties.

Product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

Compaq, Deskpro, Fastart, Compaq Insight Manager, LTE, PageMarq, Systempro, Systempro/LT, ProLiant, LicensePaq, QVision, SLT, ProLinea, SmartStart, NetFlex, DirectPlus, QuickFind, RemotePaq, BackPaq, TechPaq, SpeedPaq, QuickBack, PaqFax, registered United States Patent and Trademark Office.

Systempro/XL, Net1, SilentCool, LTE Elite, Presario, SmartStation, MiniStation, Vocalyst, PageMate, SoftPaq, FirstPaq, SolutionPaq, EasyPoint, EZ Help, MaxLight, MultiLock, QuickBlank, QuickLock, and UltraView are trademarks and/or service marks of Compaq Computer Corporation.

Other product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

©1995 Compaq Computer Corporation. Printed in the U.S.A.

Microsoft, Windows, Windows NT, Windows NT Advanced Server, SQL Server for Windows NT are trademarks and/or registered trademarks of Microsoft Corporation.

NetBench™ is a trademark of Ziff-Davis Publishing Company

Compaq ProLiant 4500 Servers in 1000+ User Environments
Second Edition (November 1995)
Document Number 182A/1195

Downloaded from www.Manualslib.com manuals search engine

INTRODUCTION

The primary application environments for Compaq servers included in this study were on-line transaction processing (OLTP) and file services. The results demonstrate that the Compaq ProLiant 4500 can support more than 1000 real users in demanding customer environments.

In this investigation the ProLiant 4500 was tested using industry-standard benchmarks in each application environment. Benchmarks are designed to stress the server until the server reaches the maximum potential in one of its subsystems. The benchmarks used in this study were chosen to drive the server to its maximum capacity in order to understand the upper bounds on performance.

The ProLiant 4500 was chosen for this study since it is the latest in the ProLiant product family from Compaq Systems Division. (The Compaq SystemPro XL, and the Compaq ProLiant 2000 and 4000 also included the same system architecture.) The ProLiant 4500 supports the latest in options and processors and allowed the investigators to determine which subsystem was being driven to its limit. The results obtained with the ProLiant 4500 can be used to extrapolate the results to other Compaq servers with a similar architecture, especially in a single processor environment such as NetWare.

This paper describes the benchmarks and results obtained. It also interprets the results as they relate to a real customer environment for both transaction processing and file services. In the transaction processing environment, the server was tested using TPC-C, a benchmark from the Transaction Processing Performance Council, running under the UnixWare operating system. In the file and print services environment, the server was tested using NetBench™, revision 3.0, running under the NetWare operating system, version 3.11. This paper also includes evidence from customer experience that can be used as a guideline to relate NetBench users to real users.

1000+ USERS IN AN OLTP ENVIRONMENT

A typical on-line transaction processing environment consists of a server attached to many different users via a network. These users make inquiries, changes, and updates across the network to information in a database residing on the server. Different elements of the server which might limit performance in an OLTP environment include the processor-memory subsystem, the network throughput and the Network Interface Controller (NIC), the I/O bus (such as EISA or PCI), and the disk subsystem.

Different elements of the server which might limit performance in an OLTP environment include the processor-memory subsystem, the network throughput and the network interface controller, the I/O bus (such as EISA or PCI), and the disk subsystem.

OLTP applications strain the processor-memory subsystem since the applications generate a relatively large number of cache misses. As long as the memory hierarchy can "keep up" with the processor, additional processor power will yield a higher number of transactions per unit time. (The TPC benchmarks use either transactions per second, tps, or per minute, tpm, as their metric.) The memory hierarchy performance is dictated by several factors including the size of the cache, the speed of the main memory DRAM, and the speed and congestion of the bus interconnecting the processor to the main memory.

OLTP applications strain the processor-memory subsystem since the applications generate a relatively large number of cache misses.

NICs communicate requests and responses via the network to and from the database server. The network transactions are very small by nature, both in terms of the amount of data transferred and the number of requests and replies transferred, when compared to those in file services environment. As long as a NIC capable of bus-mastering is used (so that it does not consume processor bandwidth), the performance of the NIC has a small impact on the OLTP performance.

NICs communicate requests and responses via the network to and from the database server.

The disk subsystem stores both the database and the transaction logs. The performance of the disk subsystem can affect the performance of an OLTP environment; however, most OLTP environments, especially in benchmarks, are set up with enough disks to allow for concurrent access to disks reducing latency, which minimizes negative impact on performance.

Downloaded from www.Manualslib.com manuals search engine

The I/O bus transfers data between the disk subsystem and memory and transfers requests and replies between memory and the NICs.

The I/O bus transfers data between the disk subsystem and memory and transfers requests and replies between memory and the NICs. If the I/O bus were to become fully clogged, it could limit performance in an OLTP environment. However, the amount of data transferred in an OLTP environment is relatively small, and I/O bus utilization levels remain quite low.

One advantage of the Compaq ProLiant 4500 system architecture in the transaction processing environment is that it allows concurrent access of system memory by the processors and the I/O subsystem. This allows more of the processors' performance to be realized since they are not blocked from memory access by I/O peripherals.

OLTP Benchmark Description

The industry standard benchmark for OLTP systems is the Transaction Processing Performance Council's TPC-C. Their earlier benchmarks were the TPC-A and TPC-B. TPC-A and TPC-B have been replaced by the TPC-C since it more closely represents real-world environments. (The council no longer accepts submissions of TPC-A and TPC-B results for inclusion in their reports.)

The TPC-C benchmark simulates a real user as closely as possible by building in delays to simulate user think time, user typing delays, etc.

TPC-C has a mixture of read-only and update-intensive transactions that simulate the activities found in complex, real world, OLTP application environments. It focuses on the response time for one transaction type, the highly interactive order-entry, as the overall benchmark metric. At the same time the server must be running a certain percentage of four other transaction types. By contrast, TPC-A and TPC-B have only one transaction type.

The TPC-C benchmark simulates a real user as closely as possible by building in delays to simulate user think time, user typing delays, etc. The value of the TPC-C benchmark is further increased because all benchmark results are documented and submitted to the council which audits the results. The council also publishes the results so that customers are compare them.

Results

The TPC-C, version 3, benchmark test was conducted on a Compaq ProLiant 4500 running the Novell UnixWare, release 2.0.1 and the DBMS Sybase System 10.0.2. A full disclosure report of this run is available from Compaq [ref-tpcc]. Table 1 shows the system parameters.

Table 1: System Parameters Running TPC-C Benchmark

System Components	Quantity	Description
Processor	4	100 MHz Pentium with 512K L2 cache, 2MB L3 cache per board
Memory	1	1024 MB
Disk Controllers	5	SMART SCSI Array Controller
Disks	53	2.1 GB Hot Pluggable SCSI-2 Drives
NICs	1	Compaq NetFlex-2 Ethernet

Table 2 shows the results of the testing. The Compaq ProLiant 4500 achieved a throughput of 1516.77 tpmC. It should be noted that the TPC-C, version 2, reported the throughput in transactions per second (tps) as opposed to the transactions per minute (tpm) shown here for version 3. There were 1440 simulated users in this system. The users were simulated in the benchmarking environment on four Compaq ProLiant 2000 systems running the Remote Terminal Emulation program to simulate 1440 users.

The Compaq ProLiant 4500 achieved a throughput of 1516.77 tpmC.

Downloaded from www.Manualslib.com manuals search engine

Table 2: Results of TPC-C Benchmark Test on a Compaq ProLiant 4500 5/100

TPC-C Throughput	1516.77 tpmC
Price:Performance	$319/tpmC
Number of Users	1440

There were 1440 simulated users in this system.

Interpretation of Results

Because the TPC-C benchmark simulates "real users" and places an upper limit on the response time, the results shown in Table 2 indicate that the ProLiant 4500 is able to comfortably support in excess of 1000 users in the OLTP environment.

The Compaq ProLiant 4500 running the Windows NT operating system is also able to comfortably support more than a 1000 users in the transaction processing environment.

In addition, preliminary investigations show slightly lower but comparable performance of the Sybase database with a similar set up under the Windows NT operating system. The Compaq ProLiant 4500 running the Windows NT operating system is also able to comfortably support more than a 1000 users in the transaction processing environment.

Unpublished measurements made during the audited benchmark run indicate that the processor subsystem was the limit to performance in this environment: the utilization levels of the processor-to-memory bus and of the network were within acceptable limits; utilization of the I/O bus (EISA) was measured at under 30%; and the disk subsystem was tuned to provide maximum performance at minimum cost. For these reasons, it is reasonable to expect higher performance, as indicated by the tpm, and more users as higher performance processors become available for the ProLiant 4500.

1000+ USERS IN THE FILE SERVICES ENVIRONMENT

The vast majority of PC-architecture based servers are installed in a file services environment. Historically, these installations were intended to allow for the sharing of files and printers. In time, mail services were added. Recently sharing of office automation applications across the network has been gaining in popularity. Although each of these application families places different demands on the server, they all share fundamentals of network topologies and common file access. Many networks are used for all of these applications.

In the file services environment a server is attached to one or more local area networks (LANs) through one or more network interface controllers (NICs). The LANs connect the server to the user nodes and to other servers. The server acts as the central repository for shared files: documents, mail, applications, etc. Traffic on the LAN consists of messages to and from the users and the server, and occasionally administrative messages between servers. The throughput requirement of the server in this environment varies throughout the day.

The amount of hand holding the NIC requires from the processor can directly impact the perceived performance of the server in the file services environment. If the NIC is not a bus master or if it generates many interrupts per packet, the host processor will be required to spend many cycles handling the NIC. In a single user desktop system, this may not impact overall performance. However in the heavily loaded server environment these cycles are required for other tasks.

Several factors can limit performance of handling client requests.

When a client requests a block of data from a file on the server, the server allocates some space in its memory to store the block. The block of data may be a shared file or a copy of the executable of a word processing package. The block is transferred from the disk and is stored in the area in server memory known as the disk cache. This block is then transferred from the disk cache via the network to the client. (A block to be written would travel in the reverse direction.) If the same block is requested again by any user before it is removed from the cache, the block will be

Downloaded from www.Manualslib.com manuals search engine

The final potential limiter of this environment is the physical LAN.

available in memory. Since it is already in the disk cache, the block can be accessed more quickly than if it were on the disk. This is referred to as a "hit" in the disk cache.

Several factors can limit performance of handling client requests.

- The first of these is the size of the disk cache. If it is too small, or if accesses are too random, there will be a large number of misses which must be satisfied by accessing the disk.
- Disk access times impact performance.
- If the disk system is not tuned properly, then the time required to access data from the disk may impact performance.
- Another possible performance limiter is the I/O bus over which information travels between the disks, memory, and NIC. If the I/O bus is too congested, or inefficient, performance will be limited.

New 100 Mbps Ethernet NICs, such as the Compaq NetFlex-3, are now available.

The final potential limiter of this environment is the physical LAN. Ethernet is the dominant LAN environment. The theoretical maximum bandwidth of Ethernet is 10 Mbps (megabits per second). Because of protocol overheads, etc., the practical bandwidth is substantially less in the "real world." While the Ethernet can be viewed as a bus like the EISA bus the behavior of the two is very different. All of the EISA bus bandwidth can be used effectively except for certain small overheads in transmission and arbitration. When the utilization of the Ethernet increases beyond 30%, the packets trying to access the Ethernet from different machines collide.

New 100 Mbps Ethernet NICs, such as the Compaq NetFlex-3, are now available. Because they require a higher grade of wiring than most Ethernet installations currently have, their use is predominately in the server. Also available are routers with a few 100 Mbps ports that can connect to the server and several 10 Mbps ports that can connect to the clients. In 10 Mbps Ethernet, about 1.2 Megabytes per second, the wire is often the limiter of performance and how many users can supported.

Today's 100 Mbps technologies are moving the bottleneck to other parts of the system. However, with the networking subsystem capable of such high throughputs, it is feasible to consider supporting a large number of users on a single server by consolidating several servers into one.

Today's 100 Mbps technologies are moving the bottleneck to other parts of the system. However, with the networking subsystem capable of such high throughputs, it is feasible to consider supporting a large number of users on a single server by consolidating several servers into one.

FILE SERVICES BENCHMARK DESCRIPTION

One of the industry standard benchmarks for file services is NetBench published by Ziff-Davis.

One of the industry standard benchmarks for file services is NetBench published by Ziff-Davis. The tests for this investigation are based on NetBench version 3.0. The NetBench program runs on the client workstations. NetBench requires no special server software apart from the file services of the Network Operating System (NOS). Since a majority of Compaq customers use Novell's NetWare as their NOS in the file services environment, the results given here are based on NetWare, version 4.1. The NetBench console is a controlling workstation and ran under Windows 3.1. The NetBench clients were run under DOS. Each NetBench client read or wrote a block of data on the server. When the request is complete and the client has either received the data requested or has received an acknowledgment of the write, the client generates another request. The clients continue making these requests for the duration of the benchmark.

The NetBench clients make simultaneous access to the server. This is uncharacteristic of real users. Users working on workstations with disks should be able to work during at least part of the day with very few requests to the server.

The NetBench clients make simultaneous access to the server. This is uncharacteristic of real users. Users working on workstations with disks should be able to work during at least part of the day with very few requests to the server. The NetBench clients can also be set up to make back-to-back requests to the server. For example, a typical file server user will open a word processing document and then spend time editing or reading the document before saving it or opening another one. The benchmark was run in this mode during this investigation. However, this is not characteristic of a real user who would normally have a certain amount of think time between

Downloaded from www.Manualslib.com manuals search engine

requests to the server. For these and other reasons, each NetBench client in this test placed much more stress on the server than a real user.

NetBench can be run in several different configurations, each of which can stress the performance of different subsystems of the server. The NetBench NIC test concentrates on testing maximum throughput that can be sustained on the NIC cards in the server. This test caches all data in the disk cache in memory and does not involve disk accesses during the testing. This test can be configured to run at a specific client block request size. The results of this test give an upper bound on the performance the server can achieve at a specific block size granularity. This test differs from the real world in that the server disk cache hit rate is 100%. Typically, systems administrators attempt to maintain disk cache hit rates of around 85%. The performance of the disk system has essentially no impact on this benchmark.

In the NetBench Disk Mix test the server fields requests of varying block sizes and a combination of reads and writes that is defined by NetBench to simulate real usage. Each user is set up initially with a large file (minimum size of 20 MB). Each user then makes random read and write requests restricted only to its file. While real users share files, this effect is not simulated in the NetBench Disk Mix test. Therefore the disk cache hit ratio will be lower when running NetBench Disk Mix test with a large number of users than the expected disk cache hit ratio in the real world. The results of this test give a lower bound on server performance.

Test Setup

The ProLiant 4500 server was tested at the Novell Superlab. The Superlab provides access to 1000 workstations. During the period of the investigation only 903 of the clients were operational.

The ProLiant 4500 server was tested at the Novell Superlab. The Superlab provides access to 1000 workstations. During the period of the investigation only 903 of the clients were operational. The overall test set up is shown in Figure 1. The clients were connected to 10 Mbit Ethernet segments. There were up to 36 workstations on each of 28 segments. Two Sysconnect 2800 switching hubs were used. Each switching hub had fourteen 10 Mbit ports connected to client segments and one 100 Mbit port connected to the ProLiant 4500. The ProLiant 4500 had two NetFlex-3 100 BaseTX NIC cards installed.

Downloaded from www.Manualslib.com manuals search engine

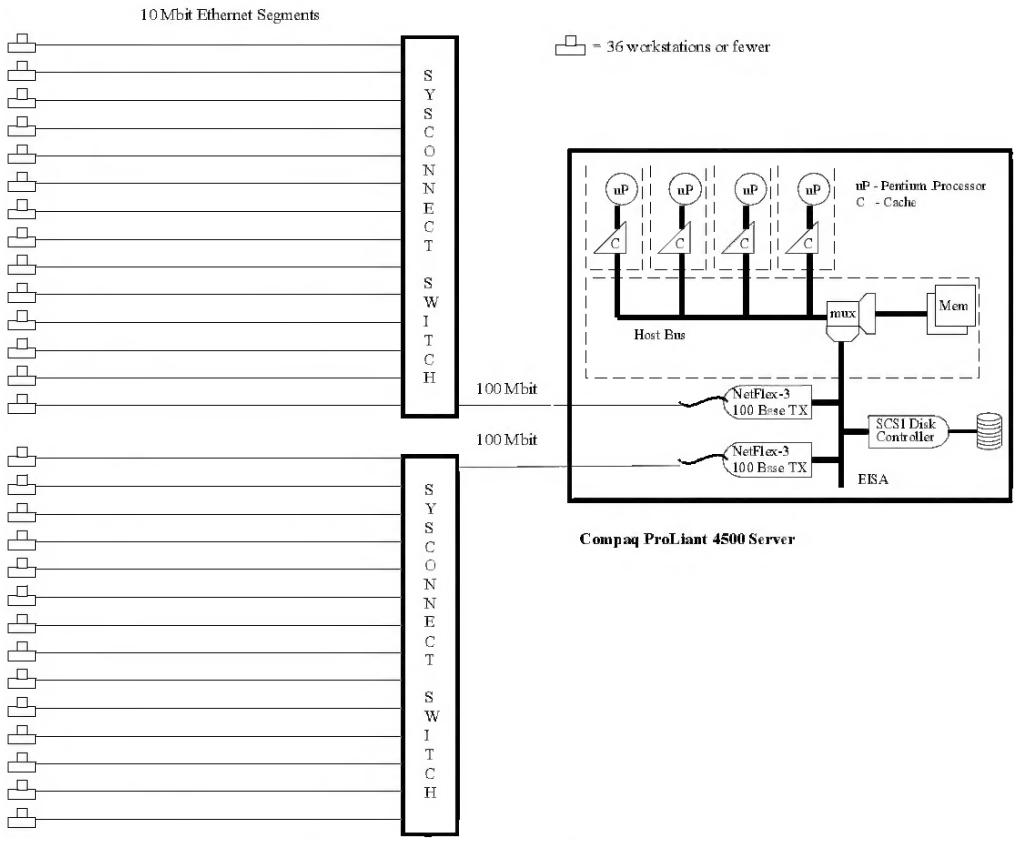

Figure 1: System setup for 1000 client NetBench run on a Compaq ProLiant 4500

Results

This test ran at the maximum stress level (no think time between iterations).

The results of the test are shown in Figure 2. The NetBench NIC test was run with a request size of 4KB. The NIC test was repeated for varying number of clients. During each run, the number of clients in each Ethernet segment was equal, balancing the client load across all segments. NetBench allows varying the stress level placed on the server by varying the think time of each client (the time that the client pauses between iterations). This test ran at the maximum stress level (no think time between iterations). Refer to the appendix for the system set up. Figure 2 shows the system throughput and the corresponding processor and EISA bus utilization when different numbers of clients were active. The following section describes the interpretation of the numbers presented.

Downloaded from www.Manualslib.com manuals search engine

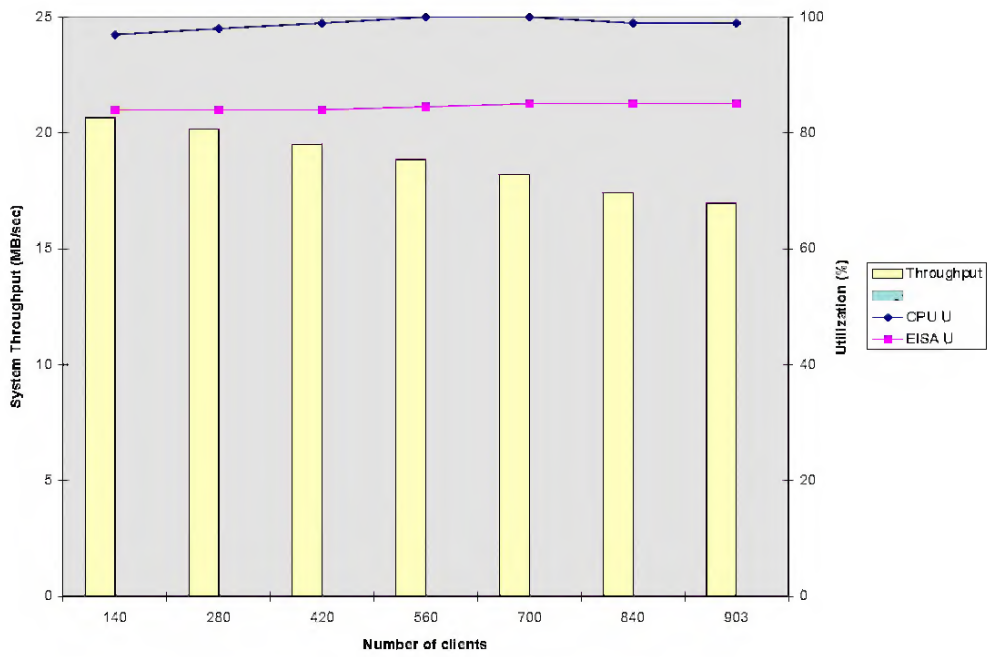

Figure 2: Throughput, processor Utilization, and EISA Utilization for the Compaq ProLiant 4500 running NetBench with 903 clients. Processor: Pentium/133 MHz; Memory: 1GB; Disk: 48 disks

Interpretation of Results

Figure 2 shows the system throughput, processor utilization and the EISA bus utilization plotted against the number of clients for the NetBench NIC test on a Compaq ProLiant 4500. As seen in the figure the throughput from the server is a maximum at 140 clients and then it drops slightly as the number of users is increased to 903. This is due to collisions in the client Ethernet segment. There were approximately 36 clients per Ethernet segment in the maximum client configuration. The clients were set to run at the maximum stress level. This means that the clients were requesting more packets from the server as soon as they received a reply for a previous request. When a larger number of very active clients are connected to the same Ethernet segment, collisions in the segment cause network inefficiencies due to the collision detect and back-off protocol used in Ethernet. Therefore, the maximum possible bandwidth in the system cannot be utilized effectively. The results presented here are for up to 903 simultaneous users.

The results presented here are for up to 903 simultaneous users.

As mentioned, the NetBench clients in these tests did not pause between getting a reply and sending out the next request. A real user would have a certain amount of think time to process the reply received before sending out another request. Also, in the real world not all users would issue requests to the server simultaneously. Therefore these NetBench throughput results can safely be extrapolated to a larger number of real users.

As mentioned, the NetBench clients in these tests did not pause between getting a reply and sending out the next request. A real user would have a certain amount of think time to process the reply received before sending out another request.

Figure 2 also shows the processor utilization and the utilization of the EISA bus. The processor utilization was measured using a NetWare Loadable Module (NLM). As can be seen, the NLM reports that the processor is almost fully utilized. In this test, the processor is processing request packets, building data packets, servicing interrupts, and initiating bus master transfers as quickly as it can; thus, the network is not the bottleneck. This high processor utilization is different than that previously observed with 10 Mbps Ethernet which runs at one-tenth the speed. It is consistent, however, because in those previous systems the network was the bottleneck and in the new systems it no longer is.

Downloaded from www.Manualslib.com manuals search engine

The EISA bus has a relatively constant utilization of 85% during the benchmark. While 85% is one of the highest EISA utilizations reported during a benchmark, it is not the performance limiter. The EISA bus protocol has been designed to have a very low arbitration and switching overhead so most of the "utilization" is used for useful work. Due to its design characteristics, it is possible to drive EISA utilization to well above 90% and still accomplish useful work if intelligent peripherals are being used.

The EISA bus has a relatively constant utilization of 85% during the benchmark. While 85% is one of the highest EISA utilizations reported during a benchmark, it is not the performance limiter. The EISA bus protocol has been designed to have a very low arbitration and switching overhead so most of the "utilization" is used for useful work. Due to its design characteristics, it is possible to drive EISA utilization to well above 90% and still accomplish useful work if intelligent peripherals are being used. The NetFlex-3 controllers and their control software have been designed to make efficient use of the bus bandwidth. As a side note, since the bottleneck in the system is in other subsystems and not in the I/O bus, this particular set up will have equal throughput on a server that has similar processor and memory architecture but has a PCI bus as the I/O bus.

Bus Utilization Numbers

While the Ethernet can be viewed as a bus like the EISA bus the behavior of the two is very different. All of the EISA bus bandwidth can be used effectively except for certain small overheads in transmission and arbitration. In this particular situation the overhead is between 10-15% of the EISA bus. The remaining bandwidth is used for data transfer. By contrast, shared Ethernet can only use about 30% of its bandwidth effectively. When the utilization of the Ethernet increases beyond 30%, the packets trying to access the Ethernet from different machines collide. When a collision is detected, the two machines back off for a random amount of time before retrying. If the same packet has multiple collisions, each time it backs off for an exponentially increasing amount of time. Therefore, with increasing collisions, the effective throughput will drop exponentially. The EISA bus has a central arbitration that controls access of all masters to the bus. There is no need to detect collisions, since only one master can use the bus at any given time. Therefore, the bandwidth is not wasted in retransmissions and silent periods (when no master is using the bus because they have to wait for the back-off timer to expire).

The maximum throughput from the system with two 100 Mbit NICs installed and 140 NetBench clients was 21 MB/sec.

The maximum throughput from the system with two 100 Mbit NICs installed and 140 NetBench clients was 21 MB/sec. Because this is a NetBench NIC test, it does not include disk access. If, in the real user environment, there are some disk accesses, part of the EISA bus bandwidth and the processor bandwidth would be required to service the disk accesses. This would reduce the total server bandwidth available for network services.

Also, the test results shown in Figure 2 are for a block request size of 4KB. In a typical user environment, 4KB is likely to be a median block size; however, it is not likely to be an average block size. (Note: Although the requested block size is 4KB, the largest packet which can be transferred on Ethernet is 1514 bytes, so the block is sent in several packets.) With larger block size requests, the network and bandwidth within the server are used more efficiently. The theoretical maximum throughput achievable through the two NICs is 50 MB/sec (400 Mbit/sec, full duplex). The NetBench NIC test traffic is primarily in one direction. since the clients do not write to the server, they only read from it. In the NetBench NIC test the theoretical maximum throughput achievable is 25 MB/sec (200 Mbit/sec), or one-half the full duplex case. Figure 2 shows that the peak system throughput achieved (as reported by NetBench) was 21 MB/sec. The remaining bandwidth of the network is used to transfer the requests, packet headers, acknowledgments, and so forth. When the packet size is reduced to 512 bytes, the peak system throughput will be reduced, perhaps to as low as 15 MB/sec. Therefore if the users are expected to transfer mostly small packets, the maximum available bandwidth will be reduced.

It is important to be able to relate the results of the NetBench benchmarks to the behavior of "real" NetWare clients.

Extrapolation of NetBench Results to Typical Usage

It is important to be able to relate the results of the NetBench benchmarks to the behavior of "real" NetWare clients. The NetBench clients used in this investigation were set up to stress the server as much as possible. This allows one to fully load a server with a handful of workstations in the test environment. The NetBench clients are configured this way since it is not always

Downloaded from www.Manualslib.com manuals search engine

possible to have a large number of clients available for tests. It is also difficult to instrument servers in a "live" user network: network administrators typically do not want to allow wires to be soldered to their production servers for the purpose of "vendor research" .

Since the NetBench client was configured to issue repeated back-to-back requests for blocks of data, the NetBench client can be viewed as a "worst-case" user. Real users require some think time, and it is highly unlikely that on a network of 1000 users all will issue simultaneous requests for large amounts of data.

Since the NetBench client was configured to issue repeated back-to-back requests for blocks of data, the NetBench client can be viewed as a "worst-case" user. Real users require some think time, and it is highly unlikely that on a network of 1000 users all will issue simultaneous requests for large amounts of data.

The throughput required by a client in the real world varies from minute to minute and from user to user. The throughput depends on the nature of the work being performed. If a network administrator is aware of or is able to predict the throughput requirement of a typical user in the administrators installation, then the number of clients that can be supported by a server can be derived by dividing the maximum throughput available from the server (around 20 MB/sec in this case) by the throughput requirement per client.

If 1000 users were supported on the server in this test, and they all accessed the server simultaneously, then the throughput available per user would be as shown in Table 3 if the users are not expected to access the server simultaneously, then the throughput available per client increases.

Table 3: Throughput available from a Compaq ProLiant 4500 for a single client

Total server throughput	Number of clients	Throughput available per client
Low Contention: 21 MB/sec	1000	21 KB/sec
High Contention: 18 MB/sec	1000	18 KB/sec

CONCLUSION

This paper shows the capability of the ProLiant 4500 in the OLTP environment and the file services environment and has shown that the server has the capability to support 1000+ users adequately in these environments.

This paper shows the capability of the ProLiant 4500 in the OLTP environment and the file services environment and has shown that the server has the capability to support 1000+ users adequately in these environments. In the transaction processing environment the server was able to support 1440 TPC-C clients, which is equivalent to supporting 1440 real users while the server is doing certain processing in the background. In the file services environment, the maximum available throughput from the server is over 20MB/sec.

One Compaq customer considering consolidation of their servers estimates that approximately 1000 users will require 5MB/sec throughput from the server. This number was derived from the typical network utilization on each sub-net times the number of sub-nets. This number was estimated for a predominately office automation environment, with some order entry functions. Even in the worst case, the ProLiant 4500 or any other Compaq server with the same system architecture machine will be able to handle that throughput easily.

However, having 1000 or more users on a single machine is really more than an issue of throughput if -- 1000 users are depending on a single machine and that machine fails, then there are 1000 people demanding its immediate return. Compaq has realized this from the dawn of the PC architecture based server, and for that reason has built availability features into its product offerings:

- RAID controllers for storage (introduced in 1989)
- Hot plug disk drives (introduced in 1992)
- Pre-failure warranty monitoring of key components (introduced in 1992)

Downloaded from www.Manualslib.com manuals search engine

- Fully integrated UPS support (introduced in 1994)
- Redundant power supplies (introduced in 1994)
- Standby Recovery Server support (introduced in 1995)
- On Line backup processors
- Advanced ECC memory
- Compaq Insight Manager, with failure paging
- 7x24 technical support hot lines
- world-wide service and support offerings
- driver integration and verification
- SmartStart configuration helper

And the list continues to grow with every generation.

The bottom line is that Compaq provides the hardware that supports 1000+ users and that network administrators don't mind supporting them on.

The bottom line is that Compaq provides the hardware that supports 1000+ users and that network administrators don't mind supporting.

REFERENCE

[ref-tpcc] TPC Benchmark C Full Disclosure Report for ProLiant 4500 5/100 Model-1 using UnixWare 2.0.2 and Sybase System 10.0.2, *Compaq Computer Corporation,* First Edition, May 1995.

Downloaded from www.Manualslib.com manuals search engine